Equity and Equality in Fair Federated Learning

Hamid Mozaffari, Amir Houmansadr

College of Information and Computer Sciences, University of Massachusetts Amherst

Introduction and Motivation

- **Federated Learning (FL)** enables data owners to train a shared global model without sharing their private data.
- Unfortunately, FL is susceptible to an intrinsic **fairness issue**:
  - Due to heterogeneity in clients' data distributions, the final trained model can give disproportionate advantages across the participating clients.
- In this work, we look at FL fairness with **two different lenses**:
  - ***Equality***: whose goal is providing similar performances for all individual clients.
  - ***Equity***: whose goal is providing similar performances across all groups of clients (i.e., groups of majority and minority), where a group is defined as a set of clients with similar data distributions.

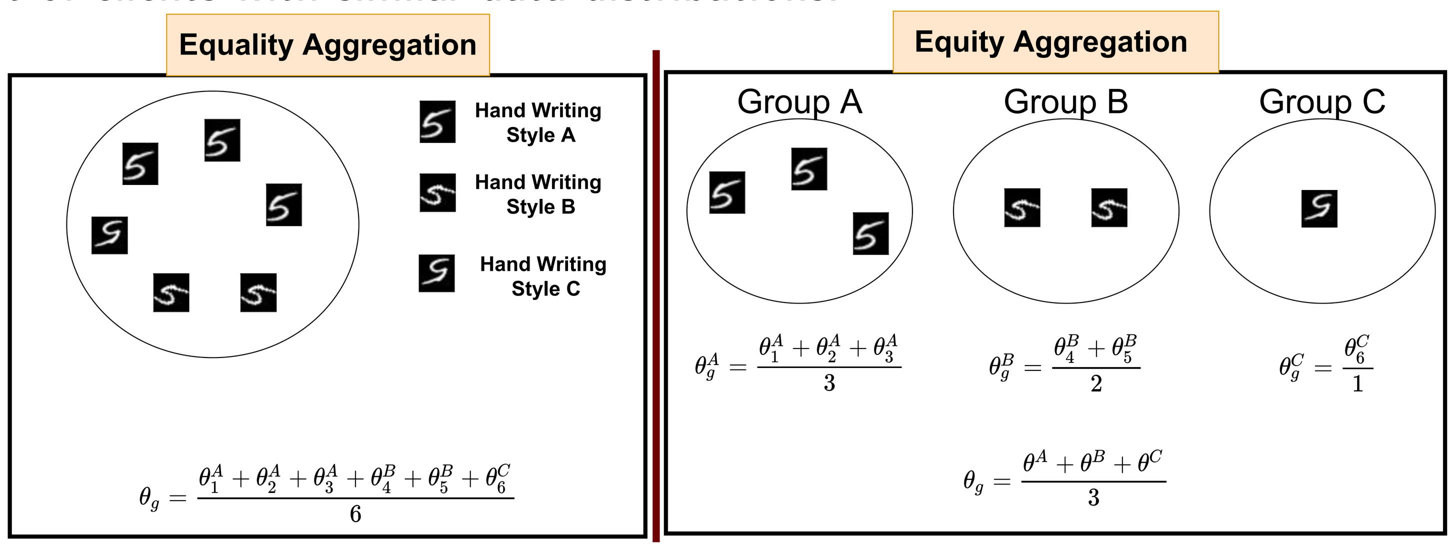

- The key question we try to answer is: *Can we design an efficient federated learning algorithm that achieves both equality and equity concurrently?*

E2FL: Equal and Equitable Federated Learning

- The key insight used in E2FL is converting the problem of model weight optimization (in standard FL) to the problem of **ranking** model edges.
- A single E2FL round with six clients from three groups and network of 6 edges.
  - Server: Initialization phase (only for round $t = 1$)
  - Clients: Calculating the **ranks** (for each round t).
  - Server: **Majority vote** (for each round t).

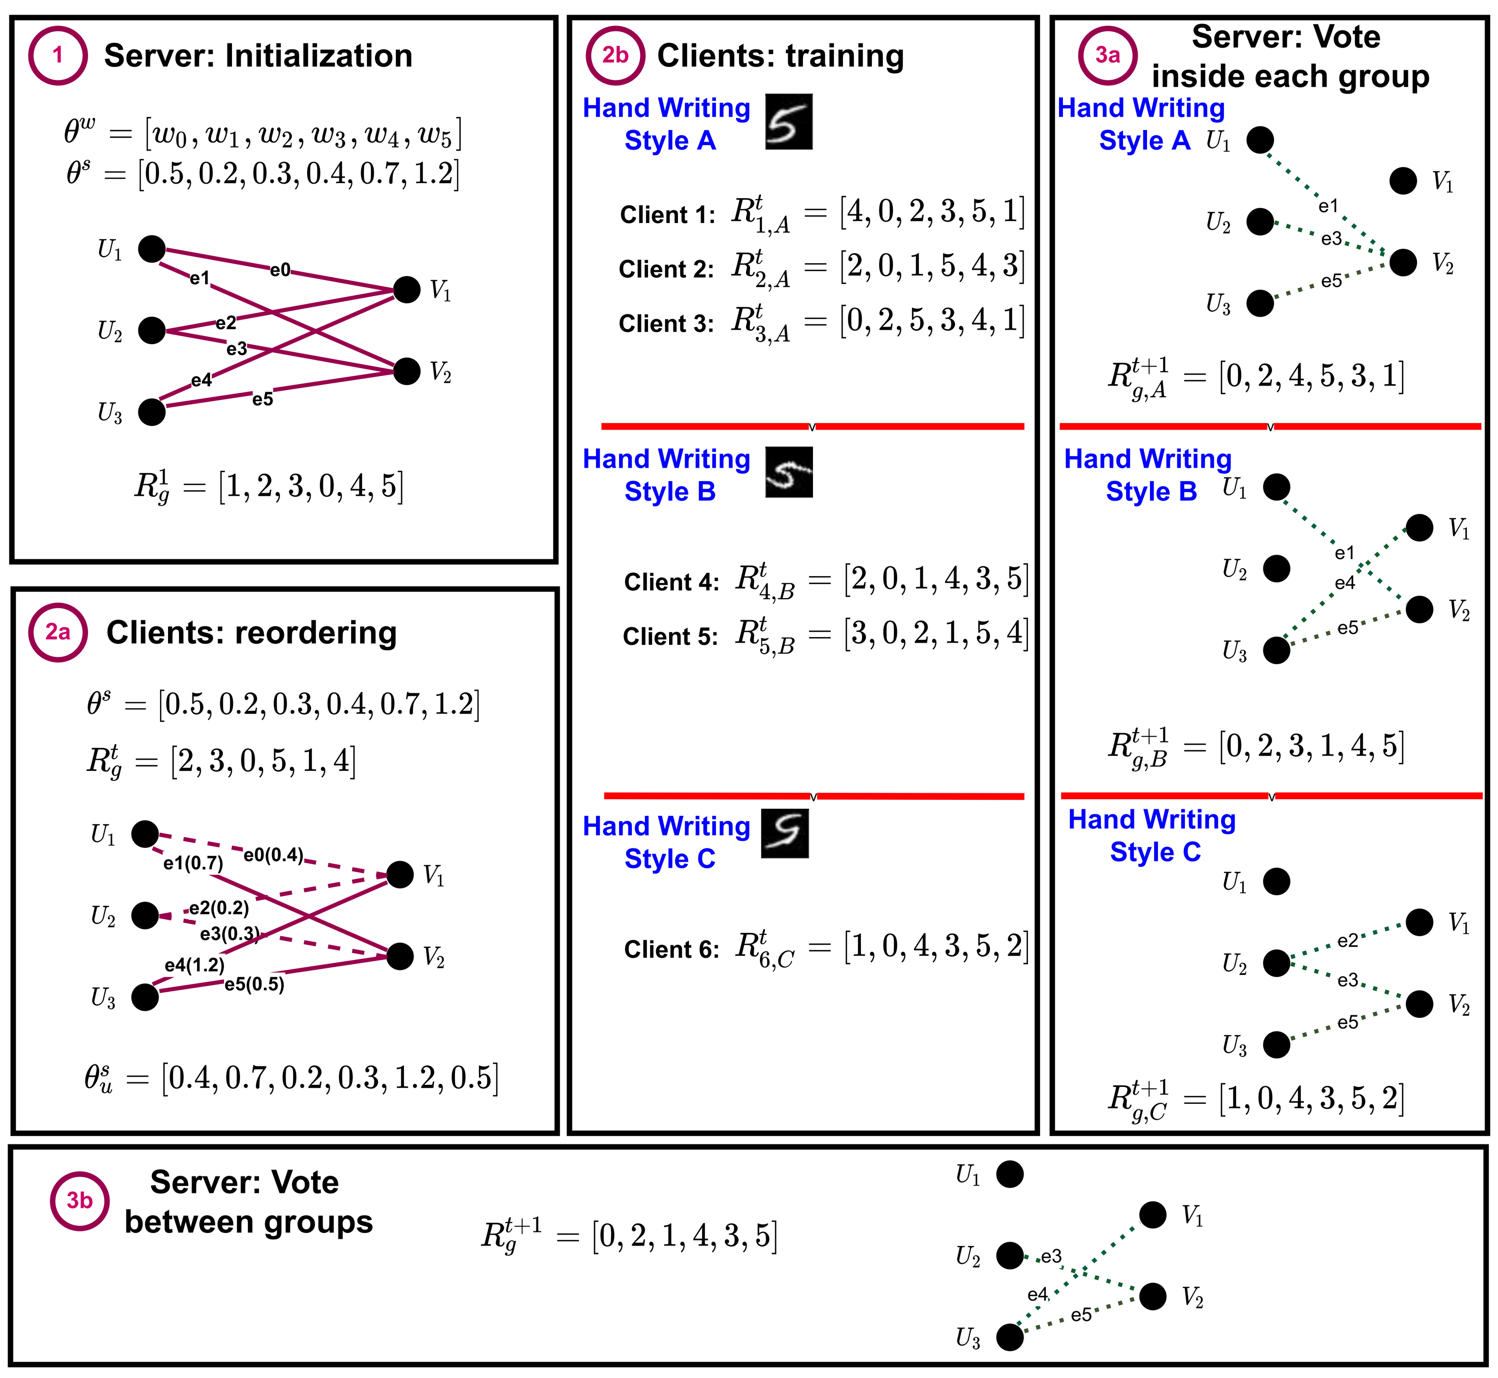

E2FL when Group IDs are Unknown

- In this setting, there are federated clients that have small amount of data with **no known protected attribute**.
- For instance, people with their own style of handwriting want to learn a global model without sharing their images of their handwriting.
- Proposed approaches for estimating Group ID:
  - Server-side: Rank **Clustering**.
  - Client-side: Lowest Loss.
  - Client-side: Entropy of the output (**OneShot Inference**, Binary Search).

Experiments: FairMNISTPerm

FairMNISTPerm: To measure the equity and equality, we release a new dataset for fairness experiments in FL application.

- We create this dataset by **rotating** the images in MNIST for each group.
- In this dataset, we assign same number of data samples to **1000 clients** with **10 different data distributions** (with different number of users in each group).
- There are **majority groups** with large number of clients, e.g., G6 with 257 clients, and there are **minority groups** with small number of clients, e.g., G1 with 8 clients.

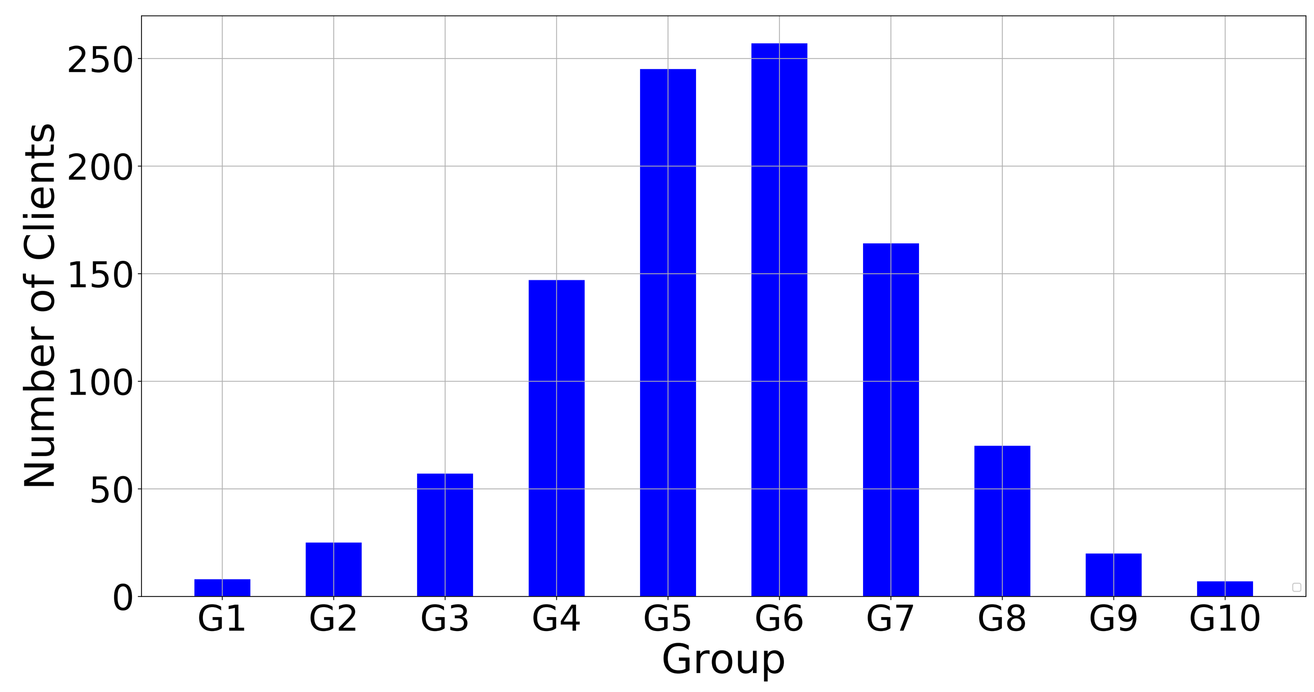

Our Results

Our experimental results on FairMNISTPerm (with 1000 clients) show that:

- Clients have **motivation** to participate in FL compared to train local models.
- **FedAvg** gives more attention to majority groups.
- **q-FFL** improves equality while worsens equity. q-FFL is helping the majority groups by ignoring the minorities.
  - q-FFL reduces the variance of accuracies for all the clients by 4% while it increases the variance between groups by 81% compared to FedAvg.
- Training 10 different FLs (i.e., **IFCA**) is not the best situation for the minorities.
- E2FL is providing equality and equity.
  - our algorithm can reduce both variance of clients and groups by 93% and 95% respectively compared to FedAvg.

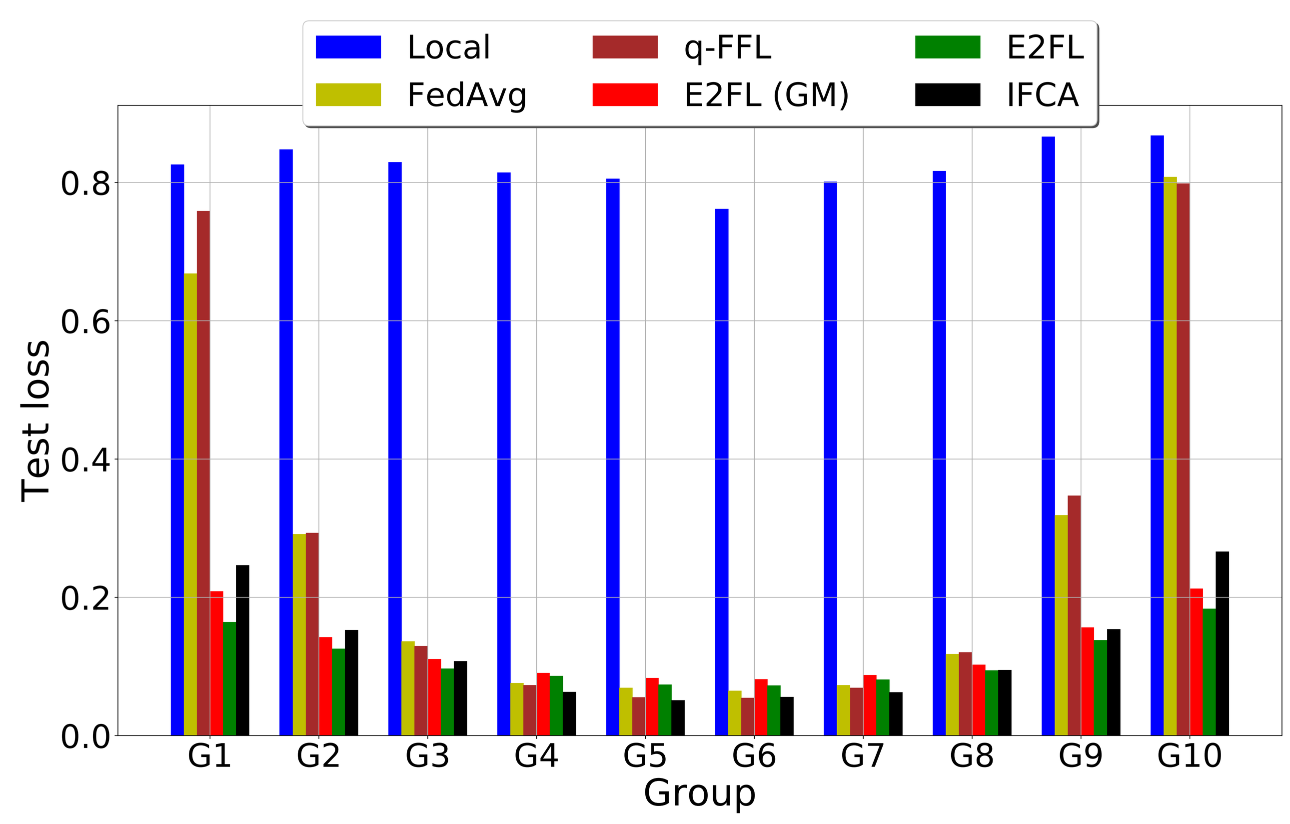

Approach	Metric									
	Group-level Fairness (Equity)				User-level Fairness (Equality)				Comm Cost	
	Avg	Worst (10%)	Best (10%)	Variance	Avg	Worst (10%)	Best (10%)	Variance	Up (MB)	Down (MB)
Local training	84.78	84.28	85.35	0.11	85.03	81.36	87.78	3.44	0	0
FedAvg	93.89	81.88	98.32	31.81	97.61	92.87	98.32	4.49	6.20	6.20
IFCA	**97.78**	**95.01**	**99.08**	1.98	**98.79**	**96.86**	**99.08**	0.35	6.20	62.0
q-FFL	92.23	77.31	98.42	57.46	97.56	93.33	98.42	4.32	6.20	6.20
Our E2FL (GM)	95.41	91.68	97.18	3.38	96.67	94.93	97.18	0.57	4.05	5.99
Our E2FL	96.52	93.90	97.81	**1.63**	97.48	96.07	97.81	**0.33**	**4.05**	**5.99**